Avionics, Navigation, and Instrumentation

The Space Shuttle faced many vehicle control challenges during ascent, as did the Orbiter during on-orbit and descent operations. Such challenges required innovations such as fly-by-wire, computer redundancy for robust systems, open-loop main engine control, and navigational aides. These tools and concepts led to groundbreaking technologies that are being used today in other space programs and will be used in future space programs. Other government agencies as well as commercial and academic institutions also use these analysis tools. NASA faced a major challenge in the development of instruments for the Space Shuttle Main Engines—engines that operated at speeds, pressures, vibrations, and temperatures that were unprecedented at the time. NASA developed unique instruments and software supporting shuttle navigation and flight inspections. In addition, the general purpose computer used on the shuttle had static random access memory, which was susceptible to memory bit errors or bit flips from cosmic rays. These bit flips presented a formidable challenge as they had the potential to be disastrous to vehicle control.

Reconfigurable Redundancy—The Novel Concept Behind the World's First Two-Fault-Tolerant Integrated Avionics System

Space Shuttle Columbia successfully concluded its first mission on April 14, 1981, with the world's first two-fault-tolerant Integrated Avionics System—a system that represented a curious dichotomy of past and future technologies. On the one hand, many of the electronics components, having been selected before 1975, were already nearing technical obsolescence. On the other hand, it used what were then-emerging technologies; e.g., time-domain-multiplexed data buses, fly-by-wire flight control, and digital autopilots for aircraft, which provided a level of functionality and reliability at least a decade ahead of the avionics in either military or commercial aircraft. Beyond the technological "nuts and bolts" of the on-board system, two fundamental yet innovative precepts enabled and shaped the actual implementation of the avionics system. These precepts included the following:

STS-1 launch (1981) from Kennedy Space Center, Florida. First crewed launch using two-fault-tolerant Integrated Avionics System.

- The entire suite of avionics functions, generally referred to as "subsystems"—data processing (hardware and software), navigation, flight control, displays and controls, communications and tracking, and electrical power distribution and control—would be programmatically and technically managed as an integrated set of subsystems. Given that new and unique types of complex hardware and software had to be developed and certified, it is difficult to overstate the role that approach played in keeping those activities on course and on schedule toward a common goal.
- A digital data processing subsystem comprised of redundant central processor units plus companion input/output units, resident software, digital data buses, and numerous remote bus terminal units would function as the core subsystem to interconnect all avionics subsystems. It also provided the means for the crew and ground to access all vehicle systems (i.e., avionics and non-avionics systems). There were exceptions to this, such as the landing gear, which was lowered by the crew via direct hardwired switches.

Avionics System Patterned After Apollo; Features and Capabilities Unlike Any Other in the Industry

The preceding tenets were very much influenced by NASA's experience with the successful Apollo primary navigation, guidance, and control system. The Apollo-type guidance computer, with additional specialized input/output hardware, an inertial reference unit, a digital autopilot, fly-by-wire thruster control, and an alphanumeric keyboard/display unit represented a nonredundant subset of critical functions for shuttle avionics to perform. The proposed shuttle avionics represented a challenge for two principal reasons: an extensive redundancy scheme and a reliance on new technologies.

Shuttle avionics required the development of an overarching and extensive redundancy management scheme for the entire integrated avionics system, which met the shuttle requirement that the avionics system be "fail operational/fail safe"—i.e., two-fault tolerant with reaction times capable of maintaining safe computerized flight control in a vehicle traveling at more than 10 times the speed of high-performance military aircraft.

Shuttle avionics would also rely on new technologies—i.e., time-domain data buses, digital fly-by-wire flight control, digital autopilots for aircraft, and a sophisticated software operating system that had very limited application in the aerospace industry of that time, even for noncritical applications, much less for "man-rated" usage. Simply put, no textbooks were available to guide the design, development, and flight certification of those technologies and only a modicum of off-the-shelf equipment was directly applicable.

Why Fail Operational/Fail Safe?

Previous crewed spacecraft were designed to be fail safe, meaning that after the first failure of a critical component, the crew would abort the mission by manually disabling the primary system and switching over to a backup system that had only the minimum capability to return the vehicle safely home. Since the shuttle's basic mission was to take humans and payloads safely to and from orbit, the fail-operational requirement was intended to ensure a high probability of mission success by avoiding costly, early termination of missions.

Early conceptual studies of a shuttle-type vehicle indicated that vehicle atmospheric flight control required full-time computerized stability augmentation. Studies also indicated that in some atmospheric flight regimes, the time required for a manual switchover could result in loss of vehicle. Thus, fail operational actually meant that the avionics had to be capable of "graceful degradation" such that the first failure of a critical component did not compromise the avionic system's capability to maintain vehicle stability in any flight regime.

The graceful degradation requirement (derived from the fail-operational/fail-safe requirement) immediately provided an answer to how many redundant computers would be necessary. Since the computers were the only certain way to ensure timely graceful degradation—i.e., automatic detection and isolation of an errant computer—some type of computerized majority-vote technique involving a minimum of three computers would be required to retain operational status and continue the mission after one computer failure. Thus, four computers were required to meet the fail-operational/fail-safe requirement. That level of redundancy applied only to the computers. Triple redundancy was deemed sufficient for other components to satisfy the fail-operational/fail-safe requirement.

Central Processor Units Were Available Off the Shelf—Remaining Hardware and Software Would Need to be Developed

The next steps included: selecting computer hardware that was for military use yet commercially available; choosing the actual configuration, or architecture, of the computer(s), data bus network, and bus terminal units; and then developing the unique hardware and software to implement the world's first two-fault-tolerant avionics.

In 1973, only two off-the-shelf computers available for military aircraft offered the computational capability for the shuttle. Both computers were basic processor units—termed "central processor units"—with only minimal input/output functionality. NASA selected a vendor to provide the central processor units plus new companion input/output processors that would be developed to specifications provided by architecture designers. At the time, no proven best practices existed for interconnecting multiple computers, data buses, and bus terminal units beyond the basic active/standby manual switchover schemes.

The architectural concept figured heavily in the design requirements for the input/output processor and two other new types of hardware "boxes" as

Interconnections Were Key to Avionics Systems Success

Shuttle Systems Elements

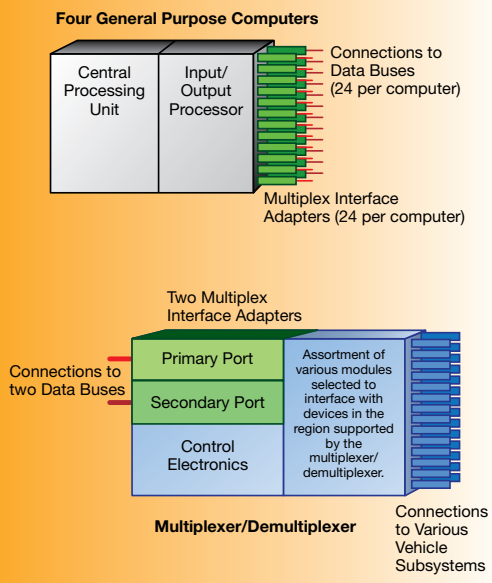

Shuttle Systems Redundancy

Diagram illustrates the eight "flight-critical" buses of the 24 buses on the Orbiter.

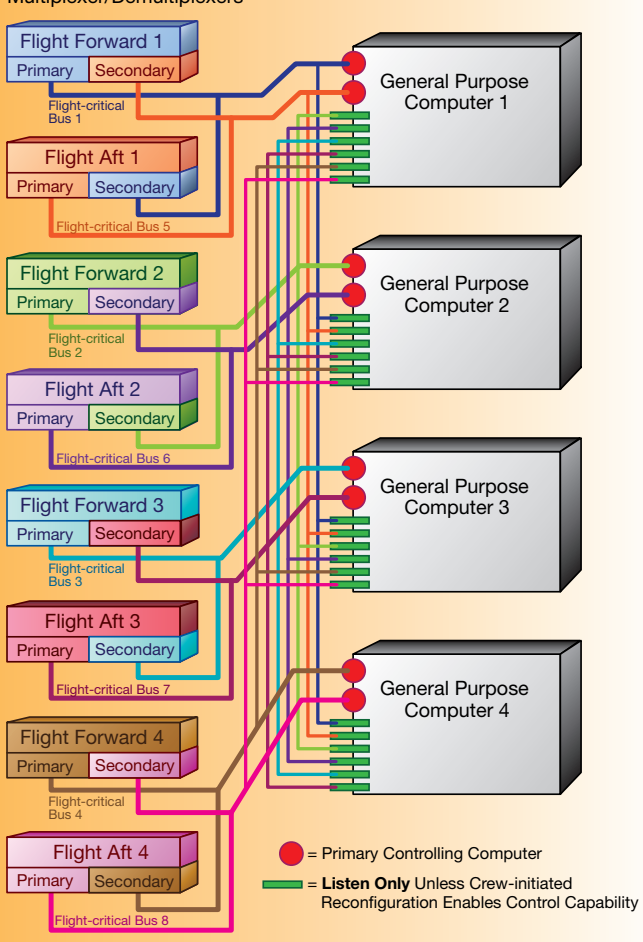

Architecture designers for the shuttle avionics system had three goals: provide interconnections between the four computers to support a synchronization scheme; provide each computer access to every data bus; and ensure that the multiplexer/demultiplexers were sufficiently robust to preclude a single internal failure from preventing computer access to the systems connected to that multiplexer/demultiplexer.

To meet those goals, engineers designed the input/output processor to interface with all 24 data buses necessary to cover the shuttle. Likewise, each multiplexer/demultiplexer would have internal redundancy in the form of two independent ports for connections to two data buses. The digital data processing subsystem possessed eight flight-critical data buses and the eight flight-critical multiplexer/demultiplexers. They were essential to the reconfiguration capability. The total complement of such hardware on the vehicle consisted of 24 data buses, 19 multiplexer/demultiplexers, and an almost equal number of other types of specialized bus terminal units.

well as the operating system software, all four of which had to be uniquely developed for the shuttle digital data processing subsystem. Each of those four development activities would eventually result in products that established new limits for the so-called "state of the art" in both hardware and software for aerospace applications.

In addition to the input/output processor, the other two new devices were the data bus transmitter/receiver units—referred to as the multiplex interface adapter—and the bus terminal units, which was termed the "multiplexer/demultiplexer." NASA designated the software as the Flight Computer Operating System. The input/output processors (one paired with each central processor unit) was necessary to interface the units to the data bus network. The numerous multiplexer/demultiplexers would serve as the remote terminal units along the data buses to effectively interface all the various vehicle subsystems to the data bus network. Each central processor unit/input/output processor pair was called a general purpose computer.

The multiplexer/demultiplexer was an extraordinarily complex device that provided electronic interfaces for the myriad types of sensors and effectors associated with every system on the vehicle. The multiplex interface adaptors were placed internal to the input/output processors and the multiplexer/demultiplexers to provide actual electrical connectivity to the data buses. Multiplex interface adaptors were supplied to each manufacturer of all other specialized devices that interfaced with the serial data buses. The protocol for communication on those buses was also uniquely defined.

The central processor units later became a unique design for two reasons: within the first several months in the field, their reliability was so poor that they could not be certified for the shuttle "man-rated" application; and following the Approach and Landing Tests (1977), NASA found that the software for orbital missions exceeded the original memory capacity. The central processor units were all upgraded with a newer memory design that doubled the amount of memory. That memory flew on Space Transportation System (STS)-1 in 1981.

Although the computers were the only devices that had to be quad redundant, NASA gave some early thought to simply creating four identical strings with very limited interconnections. The space agency quickly realized, however, that the weight and volume associated with so much additional hardware would be unacceptable. Each computer needed the capability to access every data bus so the system could reconfigure and regain capability after certain failures. NASA accomplished such reconfiguration by software reassignment of data buses to different general purpose computers.

The ability to reconfigure the system and regain lost capability was a novel approach to redundancy management. Examination of a typical mission profile illustrates why NASA placed a premium on providing reconfiguration capability. Ascent and re-entry into Earth's atmosphere represented the mission phases that required automatic failure detection and isolation capabilities, while the majority of on-orbit operations did not require full redundancy when there was time to thoroughly assess the implications of any failures that occurred prior to re-entry. When a computer and a critical sensor on another string failed, the failed computer string could be reassigned via software control to a healthy computer, thereby providing a fully functional operational configuration for re-entry.

The Costs and Risks of Reconfigurable Redundancy

The benefits of interconnection flexibility came with costs, the most obvious being increased verification testing needed to certify each configuration performed as designed. Those activities resulted in a set of formally certified system reconfigurations that could be invoked at specified times during a mission. Other less-obvious costs stemmed from the need to eliminate single-point failures. Interconnections offered the potential for failures that began in one redundant element and propagated throughout the entire redundant system—termed a "single-point failure"—with catastrophic consequences. Knowing such, system designers placed considerable emphasis on identification and elimination of failure modes with the potential to become single-point failures. Before describing how NASA dealt with potential catastrophic failures, it is necessary to first describe how the redundant digital data processing subsystem was designed to function.

Establishing Synchronicity

The fundamental premise for the redundant digital data processing subsystem operation was that all four general purpose computers were executing identical software in a time-synchronized fashion such that all received the exact same data, executed the same computations, got the same results, and then sent the exact same time-synchronized commands and/or data to other subsystems.

Maintenance of synchronicity between general purpose computers was one of the truly unique features of the newly developed Flight Computer Operating System. All four general purpose computers ran in a synchronized fashion that was keyed

Shuttle Single Event Upset Environment

Five general purpose computers—the heart of the Orbiter's guidance, navigation, and flight control system—were upgraded in 1991. The iron core memory was replaced with modern static random access memory transistors, providing more memory and better performance. However, the static random access memory computer chips were susceptible to single event upsets: memory bit flips caused by high-energy nuclear particles. These single event upsets could be catastrophic to the Orbiter because general purpose computers were critical to flights since one bit flip could disable the computer.

An error detection and correction code was implemented to "fix" flipped bits in a computer word by correcting any single erroneous bit. Whenever the system experienced a memory bit flip fix, the information was downlinked to flight controllers on the ground in Houston, Texas. The event time and the Orbiter's ground track resulted in the pattern of bit flips around the Earth.

The bit flips correlated with the known space radiation environment. This phenomena had significant consequences for error detection and correction codes, which could only correct one error in a word and would be foiled by a multi-bit error. In response, system architects selected bits for each word from different chips, making it almost impossible for a single particle to upset more than one bit per word.

In all, the upgraded Orbiter general purpose computers performed flawlessly in spite of their susceptibility to ionizing radiation.

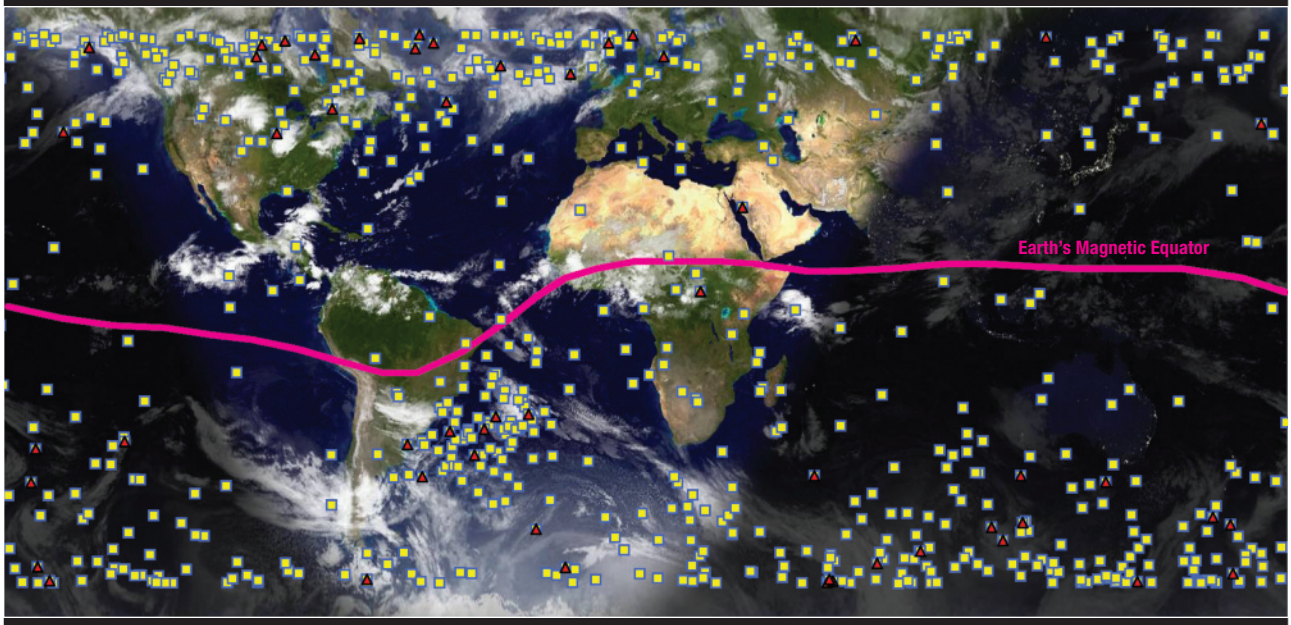

Single event upsets are indicated by yellow squares. Multi-bit single event upsets are indicated by red triangles. In these single events, anywhere from two to eight bits were typically upset by a single charged particle.

to the timing of the intervals when general purpose computers were to query the bus terminal units for data, then process that data to select the best data from redundant sensors, create commands, displays, etc., and finally output those command and status data to designated bus terminal units.

That sequence (input/process/output) repeated 25 times per second. The aerodynamic characteristics of the shuttle dictated the 25-hertz (Hz) rate. In other words, the digital autopilot had to generate stability augmentation commands at that frequency for the vehicle to retain stable flight control.

The four general purpose computers exchanged synchronization status approximately 350 times per second. The typical failure resulted in the computer halting anything resembling normal operation.

A fish-eye view of the multifunction electronic display subsystem—or "glass cockpit"—in the fixed-base Space Shuttle mission simulator at Johnson Space Center, Texas.

Early Detection of Failure

NASA designed the four general purpose computer redundant set to gracefully degrade from either four to three or from three to two members. Engineers tailored specific redundancy management algorithms for dealing with failures in other redundant subsystems based on knowledge of each subsystem's predominant failure modes and the overall effect on vehicle performance.

NASA paid considerable attention to means of detecting subtle latent failure modes that might create the potential for a simultaneous scenario. Engineers scrutinized sensors such as gyros and accelerometers in particular for null failures. During orbital operation, the vehicle typically spent the majority of time in a quiescent flight control profile such that those sensors were operating very near their null points. Prior to re-entry, the vehicle executed some designed maneuvers to purposefully exercise those devices in a manner to ensure the absence of permanent null failures. The respective design teams for the various subsystems were always challenged to strike a balance between early detection of failures vs. nuisance false alarms, which could cause the unnecessary loss of good devices.

Decreasing Probability of Pseudo-simultaneous Failures

There was one caveat regarding the capability to be two-fault tolerant—the system was incapable of coping with simultaneous failures since such failures obviously defeat the majority-voting scheme. A nuance associated with the practical meaning of "simultaneous" warranted significant attention from the designers. It was quite possible for internal circuitry in complex electronics units to fail in a manner that wasn't immediately apparent because the circuitry wasn't used in all operations. This failure could remain dormant for seconds, minutes, or even longer before normal activities created conditions requiring use of the failed devices; however, should another unrelated failure occur that created the need for use of the previously failed circuitry, the practical effect was equivalent to two simultaneous failures.

To decrease the probability of such pseudo-simultaneous failures, the general purpose computers and multiplexer/demultiplexers were designed to constantly execute cyclic background self-test operations and cease operations if internal problems were detected.

Ferreting Out Potential Single-point Failures

Engineering teams conducted design audits using a technique known as failure modes effects analysis to identify types of failures with the potential to propagate beyond the bounds of the fault-containment region in which they originated. These studies led to the conclusion that the digital data processing subsystem was susceptible to two types of hardware failures with the potential to create a catastrophic condition, termed a "nonuniversal input/output error." As the name implies, under such conditions a majority of general purpose computers may not have received the same data and the redundant set may have

diverged into a two-on-two configuration or simply collapsed into four disparate members.

Engineers designed and tested the topology, components, and data encoding of the data bus network to ensure that robust signal levels and data integrity existed throughout the network. Extensive laboratory testing confirmed, however, that the two types of failures would likely create conditions resulting in eventual loss of all four computers.

The first type of failure and the easiest to mitigate was some type of physical failure causing either an open or a short circuit in a data bus. Such a condition would create an impedance mismatch along the bus and produce classic transmission line effects; e.g., signal reflections and standing waves with the end result being unpredictable signal levels at the receivers of any given general purpose computer. The probability of such a failure was deemed to be extremely remote given the robust mechanical and electrical design as well as detailed testing of the hardware, before and after installation on the Orbiter.

The second type of problem was not so easily discounted. That problem could occur if one of the bus terminal units failed, thus generating unrequested output transmissions. Such transmissions, while originating from only one node in the network, would nevertheless propagate to each general purpose computer and disrupt the normal data bus signal levels and timing as seen by each general purpose computer. It should be mentioned that no amount of analysis or testing could eliminate the possibility of a latent, generic software error that could conceivably cause all

Loss of Two General Purpose Computers Tested Resilience

Space Shuttle Columbia (STS-9) makes a successful landing at Dryden Flight Research Center on Edwards Air Force Base runway, California, after reaching a fail-safe condition while on orbit.

Shuttle avionics never encountered any type (hardware or software) of single-point failure in nearly 3 decades of operation, and on only one occasion did it reach the fail-safe condition. That situation occurred on STS-9 (1983) and demonstrated the resiliency afforded by reconfiguration.

While on-orbit, two general purpose computers failed within several minutes of each other in what was later determined to be a highly improbable, coincidental occurrence of a latent generic hardware fault. By definition, the avionics was in a fail-safe condition and preparations were begun in preparation for re-entry into Earth's atmosphere. Upon cycling power, one of the general purpose computers remained failed while the other resumed normal operation. Still, with that machine being suspect, NASA made the decision to continue preparation for the earliest possible return. As part of the preparation, sensors such as the critical inertial measurement unit, which were originally assigned to the failed computer, were reassigned to a healthy one. Thus, re-entry occurred with a three-computer configuration and a full set of inertial measurement units, which represented a much more robust and safe configuration.

The loss of two general purpose computers over such a short period was later attributed to spacelight effects on microscopic debris inside certain electronic components. Since all general purpose computers in the inventory contained such components, NASA delayed subsequent flights until sufficient numbers of those computers could be purged of the suspect components.

four computers to fail. Thus, the program deemed that a backup computer, with software designed and developed by an independent organization, was warranted as a safeguard against that possibility.

This backup computer was an identical general purpose computer designed to "listen" to the flight data being collected by the primary system and make independent calculations that were available for crew monitoring. Only the on-board crew had the switches, which transferred control of all data buses to that computer, thereby preventing any "rogue" primary computers from "interfering" with the backup computer.

Its presence notwithstanding, the backup computer was never considered a factor in the fail-operational/fail-safe analyses of the primary avionics system, and—at the time of this publication—had never been used in that capacity during a mission.

Summary

The shuttle avionics system, which was conceived during the dawn of the digital revolution, consistently provided an exceptional level of dependability and flexibility without any modifications to either the basic architecture or the original innovative design concepts. While engineers replaced specific electronic boxes due to electronic component obsolescence or to provide improved functionality, they took great care to ensure that such replacements did not compromise the proven reliability and resiliency provided by the original design.

Development of Space Shuttle Main Engine Instrumentation

The Space Shuttle Main Engine operated at speeds and temperatures unprecedented in the history of spaceflight. How would NASA measure the engine's performance?

NASA faced a major challenge in the development of instrumentation for the main engine, which required a new generation capable of measuring—and surviving—its extreme operating pressures and temperatures. NASA not only met this challenge, the space agency led the development of such instrumentation while overcoming numerous technical hurdles.

Initial Obstacles

The original main engine instrumentation concept called for compact flange-mounted transducers with internal redundancy, high stability, and a long, maintenance-free life. Challenges presented themselves immediately, however. Few instrumentation suppliers were interested in the limited market projected for the shuttle. Moreover, early engine testing disclosed that standard designs were generally incapable of surviving the harsh environments. Although the "hot side" temperatures were within the realm of jet engines, no sort of instrumentation existed that could handle both high temperatures and cryogenic environments down to minus -253°C (-423°F). Vibration environments with high-frequency spectrums extending beyond commercially testable ranges of 2,000 hertz (Hz) experienced several hundred times the force of gravity over almost 8 hours of an engine's total planned operational exposure. For these reasons, the endurance requirements of the instrumentation constituent materials were unprecedented.

Engine considerations such as weight, concern for leakage that might be caused by mounting bosses, and overall system fault tolerance prompted the need for greater redundancy for each transducer. Existing supplier designs, where available, were single-output devices that provided no redundancy. A possible solution was to package two or more sensors within a single transducer. But this approach required special adaptation to achieve the desired small footprint and weight.

NASA considered the option of strategically placing instrumentation devices and closely coupling them to the desired stimuli source. This approach prompted an appreciation of the inherent simplicity and reliability afforded by low-level output devices. The avoidance of active electronics tended to minimize electrical, electronic, and electromechanical part vulnerability to hostile environments. Direct mounting of transducers also minimized the amount of intermediate hardware capable of producing a catastrophic system failure response. Direct mounting, however, came at a price. In some situations, it was not possible to design transducers capable of surviving the severe environments, making it necessary to off-mount the device. Pressure measurements associated with the combustion process suffered from icing or blockage issues when hardware temperatures dropped below freezing. Purging schemes to provide positive flow in pressure tubing were necessary to alleviate this condition.

Several original system mandates were later shown to be ill advised, such as an early attempt to achieve some measure of standardization through the use of bayonet-type electrical connectors. Early engine-level and laboratory testing revealed the need for threaded connectors since the instrumentation components could not be adequately shock-isolated to prevent failures induced by excessive relative connector motion. Similarly, electromagnetic interference assessments and observed deficiencies resulted in a reconsideration of the need for cable overbraiding to minimize measurement disruption.

Problems also extended to the sensing elements themselves. The lessons of material incompatibilities or deficiencies were evident in the area of resistance temperature devices and thermocouples. The need for the stability of temperature measurements led to platinum-element resistance temperature devices being baselined for all thermal measurements.

Aggressive engine performance and weight considerations also compromised the optimal sensor mountings. For example, it was not practical to include the prescribed straight section of tubing upstream from measuring devices, particularly for flow. This resulted in the improper loading of measuring devices, primarily within the propellant oxygen ducting. The catastrophic failure risks finally prompted the removal or relocation of all intrusive measuring devices downstream of the high-pressure oxygen turbopump. Finally, the deficiencies of vibration redline systems were overcome as processing hardware and algorithms matured to the point where a real-time synchronous vibration redline system could be adopted, providing a significant increase in engine reliability.

Wire Failures Prompted System Redesign

High temperature measurements continued to suffer brittle fine-element wire failures until the condition was linked to operation above the material recrystallization temperature of 525°C (977°F) where excessive grain growth would result. The STS-51F (1985) in-flight engine shutdown caused by the failure of multiple resistance temperature devices mandated a redesign to a thermocouple-based system that eliminated the wire embrittlement problem.

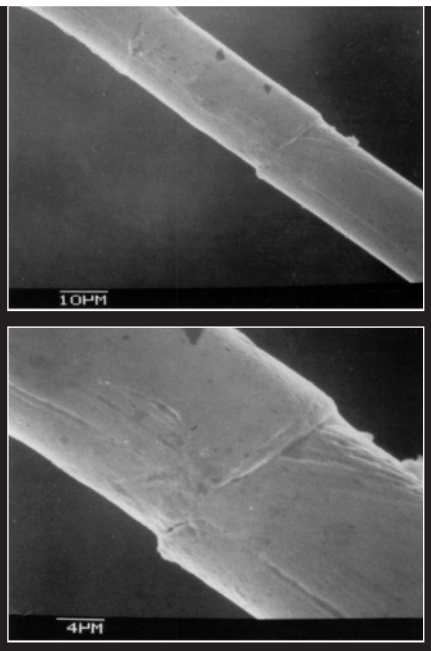

High temperatures in some engine operating environments caused fine wires used in temperature devices to become brittle, thereby leading to failures.

Weakness Detection and Solutions

In some instances, the engine environment revealed weaknesses not normally experienced in industrial or aerospace applications. Some hardware successfully passed component-level testing only to experience problems at subsystem or engine-level testing. Applied vibration spectrums mimicked test equipment limitations where frequency ranges typically did not extend beyond 2,000 Hz. The actual engine recognized no limits and continued to expose the hardware to energy above even 20,000 Hz. Therefore, a critical sensor resonance condition might only be excited during engine-level testing. Similarly, segmenting of component testing into separate vibration, thermal, and fluid testing deprived the instrumentation of experiencing the more-severe effect of combined exposures.

The shuttle's reusability revealed failure modes not normally encountered, such as those ascribed to the differences between flight and ground test environments. It was subsequently found that the microgravity exposure of each flight allowed conductive particles within instruments to migrate in a manner not experienced with units confined to terrestrial applications. Main engine pressure transducers experienced electrical shorts only during actual engine performance. During the countdown of Space Transportation System (STS)-53 (1992), a high-pressure oxidizer turbopump secondary seal measurement output pressure transducer data spike almost triggered an on-pad abort. Engineers used pressure transducers screened

by particle impact noise detection and microfocus x-ray examination on an interim basis until a hardware redesign could be qualified.

Effects of Cryogenic Exposure on Instrumentation

Cryogenic environments revealed a host of related material deficiencies. Encapsulating materials—necessary to provide structural support for fine wires within speed sensors—lacked resiliency at extreme low temperatures. The adverse effects of inadvertent exposure to liquefied gases within the shuttle's aft compartment produced functional failures due to excessively cold conditions. In April 1991, STS-37 was scrubbed when the high-pressure oxidizer turbopump secondary seal pressure measurement became erratic due to the damaging effects of cryogenic exposure of a circuit board.

Problems with cryogenics also extended to the externals of the instrumentation. Cryopumping—the condensation-driven pumping mechanism of inert gases such as nitrogen—severely compromised the ability of electrical connectors to maintain continuity. The normally inert conditions maintained within the engine system masked a problem with residual contamination of glassed resistive temperature devices used for cryogenic propellant measurements. Corrosive flux left over from the manufacturing process remained dormant for years until activated during extended exposures to the humid conditions at the launch site. STS-50 (1992) narrowly avoided a launch delay when a resistive temperature device had to be replaced just days before the scheduled launch date.

Expectations Exceeded

As the original main engine design life of 10 years was surpassed, part obsolescence and aging became a concern. Later designs used more current parts such as industry-standard electrical connectors. Some suppliers chose to invest in technology driven by the shuttle, which helped to ease the program's need for long-term part availability.

The continuing main engine ground test program offered the ability to use ongoing hot-fire testing to ensure that all flight hardware was sufficiently enveloped by older ground test units. Tracking algorithms and extensive databases permitted such comparisons.

Industry standards called for periodic recalibration of measuring devices. NASA excluded this from the Space Shuttle Main Engine Program at its inception to reduce maintenance for hardware not projected for use beyond 10 years. In practice, the hardware life was extended to the point that some engine components approached 40 years of use before the final shuttle flight. Aging studies validated the stable nature of instruments never intended to fly so long without recalibration.

Summary

While initial engine testing disclosed that instrumentation was a weak link, NASA implemented innovative and successful solutions that resulted in a suite of proven instruments capable of direct application on future rocket engines.

Unprecedented Rocket Engine Fault-Sensing System

The Space Shuttle Main Engine (SSME) was a complex system that used liquid hydrogen and liquid oxygen as its fuel and oxidizer, respectively. The engine operated at extreme levels of temperature, pressure, and turbine speed. At these levels, slight material defects could lead to high vibration in the turbomachinery. Because of the potential consequences of such conditions, NASA developed vibration monitoring as a means of monitoring engine health.

The main engine used both low- and high-pressure turbopumps for fuel and oxidizer propellants. Low-pressure turbopumps served as propellant boost pumps for the high-pressure turbopumps, which in turn delivered fuel and oxidizer at high pressures to the engine main combustion chamber.

The high-pressure pumps rotated at speeds reaching 36,000 rpm on the fuel side and 24,000 rpm on the oxidizer side. At these speeds, minor faults were exacerbated and could rapidly propagate to catastrophic engine failure.

During the main engine's 30-year ground test program, more than 40 major engine test failures occurred. High-pressure turbopumps were the source of a large percentage of these failures. Posttest analysis revealed that the vibration spectral data contained potential failure indicators in the form of discrete rotordynamic spectral signatures. These signatures were prime indicators of turbomachinery health and could potentially be used to mitigate

catastrophic engine failures if assessed at high speeds and in real time.

NASA recognized the need for a high-speed digital engine health management system. In 1996, engineers at Marshall Space Flight Center (MSFC) developed the Real Time Vibration Monitoring System and integrated the system into the main engine ground test program. The system used data from engine-mounted accelerometers to monitor pertinent spectral signatures. Spectral data were produced and assessed every 50 milliseconds to determine whether specific vibration amplitude thresholds were being violated.

NASA also needed to develop software capable of discerning a failed sensor from an actual hardware failure. MSFC engineers developed the sensor validation algorithm—a software algorithm that used a series of rules and threshold gates based on actual vibration spectral signature content to evaluate the quality of sensor data every 50 milliseconds.

Outfitted with the sensor validation algorithm and additional software, the Real Time Vibration Monitoring System could detect and diagnose pertinent indicators of imminent main engine turbomachinery failure and initiate a shutdown command within 100 milliseconds.

The Real Time Vibration Monitoring System operated successfully on more than 550 main engine ground tests with no false assessments and a 100% success rate on determining and disqualifying failed sensors from its vibration redlines. This, the first high-speed vibration redline system developed for a liquid engine rocket system, supported the main engine ground test program throughout the shuttle era.

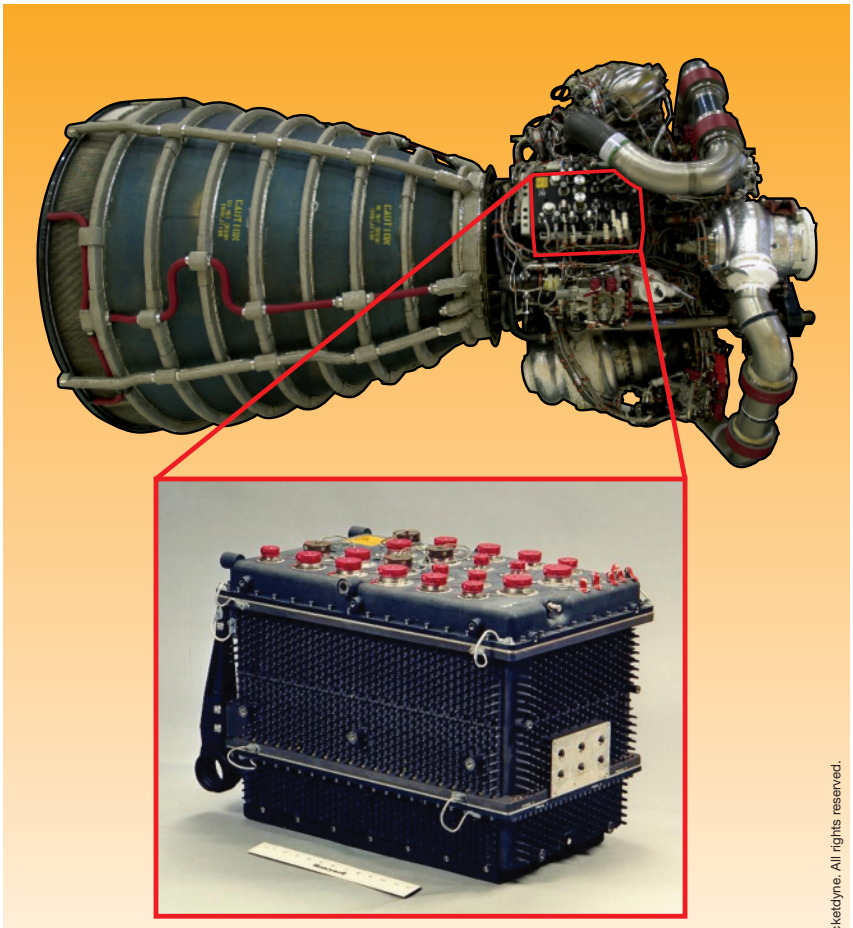

NASA's Advanced Health Monitoring System software was integrated with the Space Shuttle Main Engine controller (shown by itself and mounted on the engine) in 2007.

To prove that a vibration-based, high-speed engine health management system could be used for flight operations, NASA included a subscale version of the Real Time Vibration Monitoring System on Technology Flight Experiment 2, which flew on STS-96 (1999).

This led to the concept of the SSME Advanced Health Management System as a means of extending this protection to the main engine during ascent.

The robust software algorithms and redline logic developed and tested for the Real Time Vibration Monitoring System were directly applied to the Advanced Health Management System and incorporated into a redesigned

version of the engine controller. The Advanced Health Management System's embedded algorithms continuously monitored the high-pressure turbopump vibrations generated by rotation of the pump shafts and assessed rotordynamic performance every 50 milliseconds. The system was programmed to initiate a shutdown command in fewer than 120 milliseconds if vibration patterns indicated an instability that could lead to catastrophic failure.

The system also used the sensor-validation algorithm to monitor sensor quality and could disqualify a failed sensor from its redline suite or deactivate the redline altogether. Throughout the shuttle era, no other liquid engine rocket system in the world employed a vibration-based health management system that used discrete spectral components to verify safe operation.

Summary

The Advanced Health Management System, developed and certified by Pratt & Whitney Rocketdyne (Canoga Park, California) under contract to NASA, flew on numerous shuttle missions and continued to be active on all engines throughout the remainder of the shuttle flights.

Calibration of Navigational Aides Using Global Positioning Computers

The crew members awakened at 5:00 a.m. After 10 days in orbit, they were ready to return to Earth. By 7:45 a.m., the payload bay doors were closed and they were struggling into their flight suits to prepare for descent. The commander called for a weather report and advice on runway selection. The shuttle could be directed to any one of three landing strips depending on weather at the primary landing site. Regardless of the runway chosen, the descent was controlled by systems capable of automatically landing the Orbiter. The Orbiter commander took cues from these landing systems, controlled the descent, and dropped the landing gear to safely land the Orbiter. During their approach to the landing site, the Orbiter crew depended on a complex array of technologies, including a Tactical Air Navigation System and the Microwave Scanning Beam Landing System, to provide precision navigation. These systems were located at each designated landing site and had to be precisely calibrated to ensure a safe and smooth landing.

Touchdown Sites

Shuttle runways were strategically located around the globe to serve several purposes. After a routine mission, the landing sites included Kennedy Space Center (KSC) in Florida, Dryden Flight Research Center in California, and White Sands Test Facility in New Mexico. The transoceanic abort landing sites—intended for emergencies when the shuttle lost a main engine during ascent and could not return to KSC—were located in Zaragoza and Moron in Spain and in Istres in France. Former transoceanic abort landing sites included: Dakar, Senegal; Ben Guerir, Morocco; Banjul, The Gambia; Honolulu, Hawaii; and Anderson Air Force Base, Guam. NASA certified each site.

Error Sources

Because the ground portion of the Microwave Scanning Beam Landing and Tactical Air Navigation Systems contained moving mechanical components and depended on microwave propagation, inaccuracies could develop over time that might prove detrimental to a shuttle landing. For example, antennas could drift out of mechanical adjustment. Ground settling and external environmental factors could also affect the system's accuracy. Multipath and refraction errors could result from reflections off nearby structures, terrain changes, and day-to-day atmospheric variations.

Flight inspection data gathered by the NASA calibration team could be used to determine the source of these errors. Flight inspection involved flying an aircraft through the landing system coverage area and receiving time-tagged data from the systems under test. Those data were compared to an accurate aircraft positioning reference to determine error. Restoring integrity was easily achieved through system adjustment.

Global Positioning Satellite Position Reference for Flight Inspection

Technologies were upgraded several times since first using the Global Positioning Satellite (GPS)-enabled flight inspection system. The flight inspection system used an aircraft GPS receiver as a position reference. Differences between the system under test and the position reference were recorded, processed, and displayed in real time on board the aircraft. An aircraft position reference used for flight inspection had to be several times more accurate than the system under test. Stand-alone commercial GPS systems did not have enough accuracy for this purpose. Several techniques could be used to improve GPS positioning. Differential GPS used a ground GPS receiver installed over a known surveyed benchmark. Common mode error corrections to the GPS position were calculated and broadcast over a radio data link to the aircraft. After the received corrections were applied, the on-board GPS position accuracy was within 3 m (10 ft). A real-time accuracy within 10 cm (4 in.) was achieved by using a carrier-phase technique and tracking cycles of the L-band GPS carrier signal.

NASA built several versions of the flight inspection system customized to different aircraft platforms. Different NASA aircraft were used based on aircraft availability. These aircraft include NASA's T-39 jet (Learjet), a NASA P-3 turboprop, several C-130 aircraft, and even NASA's KC-135. Each aircraft was modified with shuttle landing system receivers and antennas. Several pallets of equipment were configured and tested to reduce the installation time on aircraft to one shift.

Summary

NASA developed unique instrumentation and software supporting the shuttle navigation aids flight inspection mission. The agency developed aircraft pallets to operate, control, process, display, and archive data from several avionics receivers. They acquired and synchronized measurements from shuttle-unique avionics and aircraft platform avionics with precision time-tagged GPS position. NASA developed data processing platforms and software algorithms to graphically display and trend landing system performance in real time. In addition, a graphical pilot's display provided the aircraft pilot with runway situational awareness and visual direction cues. The pilot's display software, integrated with the GPS reference system, resulted in a significant reduction in mission flight time.

Synergy With the Federal Aviation Administration

In early 2000, NASA and the Federal Aviation Administration (FAA) entered into a partnership for flight inspection. The FAA had existing aircraft assets to perform its mission to flight-inspect US civilian and military navigation aids. The FAA integrated NASA's carrier-phase GPS reference along with shuttle-unique avionics and software algorithms into its existing control and display computers on several flight-inspection aircraft.

The NASA/FAA partnership produced increased efficiency, increased capability, and reduced cost to the government for flight inspection of the shuttle landing aids.